Do Next Generation Networks need Path Diversity?

Luca Muscariello[1], Diego Perino[1,2], Dario Rossi[3]
[1] Orange Labs, France, first.last@orange-ftgroup.com; [2] INRIA Project Team “GANG”, France
[3] TELECOM ParisTech, France, dario.rossi@enst.fr

***Abstract*—We have currently reached a phase where big shifts in the network traffic might impose to rethink the design of current architectures, and where new technologies, being pushed into market, will act as enabler of such changes.**

Taking into account the current scenario and its likely evolution as well, in this paper we examine the case for multi-path routing within the metropolitan access network. Through an optimization framework, we undertake the analysis of several interesting aspects of the problem, such as i) the user access technology, ii) the topology of the access network and iii) the traffic locality ratio within the access.

By numerical solution of the problem we quantify the potential gain given by path-diversity: our results confirm the appeal of multi-path routing strategies both from the user and the network perspectives.

I. INTRODUCTION

The evolution of network devices, services and applications has reached a phase which imposes to rethink network design, making the case for a clean-slate design of future network paradigms, which are often referred to as “post-IP” solutions.

Indeed, network services and Internet applications keep evolving at a very fast pace. For instance, Triple Play implies that television and telephony are practically “reincarnating” into networked telco-managed services such as IPTV and VoIP. Moreover, the explosion of P2P approaches for gaming, telephony and television further blurs the distinction among data, voice and multimedia traffic. Finally, another important piece of the puzzle is constituted by the increase of traffic dynamism, that flash-crowd effects and widespread usage of application-layer overlays [1], [2], [3], [4] undoubtedly contribute to exacerbate.

Due to these new conditions, post-IP network architecture designers are looking, with increased interests towards multi-path routing. Indeed, the situation changed from the early experiences of dynamic routing, whose responsiveness to congestion has refrained its deployment because judged complex and unstable. Nowadays, a number of applications (such as conversational calls and short Web transfers) should still refrain to be split over multiple routes. At the same time, many other applications can tolerate a dynamic environment, which would definitively lead to a more efficient exploitation of the unused network capacity. This is for instance the case of rather bandwidth eager P2P applications, such as file-sharing [1], [2] and live-streaming [3], [4], that together constitute a very significant portion of the aforementioned Internet “data” traffic nowadays.

This work focuses on the design and analysis of architectures and routing mechanisms in next generation metro access networks.

The access network connects a number of equipments aggregating users with various access technologies as WiFi, radio, ADSL, FTTH running different applications. Nowadays ADSL is the main access technology while in the future mobile radio access and FTTH will become more popular enlarging the spectrum of access rates.

Let us focus for instance on the case of a metro-access network with a single gateway towards the big Internet, and define the traffic *local* when it is destined to a host *within* the access network, or *remote* when it is destined to an Internet host. In the current scenario, traffic is mostly remote (exchanges from outside the access network), as most of the services are not provided by the ISP and also a very significant fraction of P2P traffic crosses the gateway to reach faraway Internet hosts. On the one hand, it may be that increasing tendency towards topology-aware and latency-aware overlays [6], [7] will in possibly result in an increase of the traffic locality within the access network. Yet, there are other factors which may offset the potential increase of traffic locality, such as the popularity of Website like YouTube and DailyMotion, whose Internet video traffic will possibly account to more than half of all consumer traffic in the next few years [5]. If these forecasts will be fulfilled, and unless the operator have direct control over the Content Distribution Network / Service Overlay infrastructure used to distribute the video content, the access Internet gateway is likely to remain a traffic hot-spot.

As traffic and access technologies are changing , the current access ring network is expected to evolve towards more meshed topologies allowing routing to exploit path diversity.

In this paper we evaluate the potential gain offered by path diversity within the framework of optimal routing and flow control solved in a number of realistic scenarios. This model takes into account user’s access rates, and is thus useful to evaluate the performance of networks whose users have heterogeneous *access technologies*. Concerning the *traffic* description, we prefer to use an oblivious approach, in which traffic demands are not associated to a particular class of services, but are rather identified by their exogenous peak rate (i.e., the highest rate that an application would attain on unlimited capacity links).

Since the current scenario may evolve by taking rather different, but nevertheless all plausible directions, we believe that it is necessary to carefully evaluate the benefits of path-diversity under the widest possible number of scenarios.

We anticipate that our results confirms multipath to be a very appealing solution for a wide range of scenarios – i.e., for various network topologies, individual link capacities, access technology popularity and traffic locality. To the best of our knowledge this is the first attempt to promote the deployment of multipath routing on metropolitan access network based on a precise quantification of its practical benefits. Furthermore results are also valid for similar network topologies and scenarios with similar traffic characteristics.

The reminder of this paper is organized as follows. Related work are reviewed in Sec. II, while the proposed modeling framework is described in Sec. III. Then, taking into account the network and traffic assumptions detailed in Sec. IV, we quantify the benefit of multipath in Sec. V, Finally, Sec. VI concludes the paper.

II. RELATED WORK

Seminal work on optimal routing dates back to [8], [9], [10], where centralized and decentralized strategies were proposed in the very beginning of the ARPAnet project. Optimization typically explicits the problem of resource allocation under a chosen fairness criteria, which has been formulated, starting with [10], as a non linear optimization problem with linear constraints. In the last few years, intense research on multi-path routing protocols has progressed [12], [13], [14], [20], [15], [16], [17], [18], [19]. Theory of optimization has been applied to develop distributed algorithms solving a global optimization problem [12], [13], [15], [16], [17]. Control theory has been used to obtain delay stability of distributed optimal schemes [12], [14], [15], [16]. Other research considers dynamic flow level models [18], [19] in order to take into account arrivals and departures of user's sessions. The fairness criterion is realized by an end to end protocol [12], [13], [14], [20], [15], [16], [17] while can also be imposed by link scheduling [11] and flow control, in presence of congestion.

We make the assumption that multi-path routing is implemented within the network, and focus on the gain provided by path diversity in realistic scenarios, in typical access network topologies and for typical traffic demands.

To the best of our knowledge this is the first work that evaluate potential path diversity gains in the context of next generation Internet access networks with a realistic distribution of demands and access rates.

III. MULTI-PATH ROUTING FRAMEWORK

In this section we briefly describe the optimization framework, referring the reader to [10], [21] for a more detailed description. The network topology is modelled by a connected graph $G = (N, L)$, given as a set of nodes and links, with an adjacency matrix $A = [a_{ij}]$ whose elements $a_{ij} = 1$ if there exists a directional link between $i \rightarrow j$ and $a_{ij} = 0$ otherwise.

Neglecting transport layer or application layer flow semantic, we consider traffic *flows* (or *demands*) as being constituted by a traffic aggregate flowing between two access nodes connecting thousands of users. An important flow characteristic is its *exogenous rate*, that is the highest attainable rate on a link of unlimited capacity: this limit can be considered to be the maximum rate attainable by the traffic aggregate.

The network carries traffic generated by a set of demands Γ, each demand d is identified by a triple (s^d, e^d, p^d), with source $s \in \mathcal{S} \subset N$, destination $e \in \mathcal{D} \subset N$ and exogenous peak rate $p \in \mathbb{R}^+$. In our model a network flow d gets a share x_{ij}^d of the capacity c_{ij} at each link $0 \leq x_{ij}^d \leq \min(c_{ij}, p)$, with $x_{ij}^d(t)$ the fluid approximation of the rate at which the source d is sending at time t through link $i \rightarrow j$.

Route selection and bandwidth sharing are jointly modeled in the following problem, which allows traffic flows to be split among different paths available at a given node:

$$\max \sum_{i\in\mathcal{S},d\in\Gamma} U_d(\phi_i^d) - \sum_{i,j\in N} C\left(\frac{\sum_{d\in\Gamma} x_{ij}^d}{c_{ij}}\right) \tag{1}$$

subject to

$$\sum_{k\in N} a_{ik}x_{ki}^d - \sum_{j\in N} a_{ji}x_{ij}^d = \begin{cases} \phi_i^d & \text{if } \ i \in \mathcal{S} \\ -\phi_i^d & \text{if } \ i \in \mathcal{D} \\ 0 & \text{otherwise} \end{cases} \qquad \forall d \in \Gamma \tag{2}$$

$$\sum_{d\in\Gamma} x_{ij}^d \leq c_{ij} \qquad \forall i, j \in L \tag{3}$$

$$\phi_i^d \leq p^d \qquad \forall d \in \Gamma, \forall i \in \mathcal{S} \tag{4}$$

In (1), we maximize user utility $U(\cdot)$ and introduce the network cost $C(\cdot)$ that can be thought as modelling the link delay (mean delay in an M/M/1 queue), thus:

$$C(x_{ij}) = \frac{x_{ij}}{c_{ij} - x_{ij}} \tag{5}$$

$$U_d(x) = w_d(1-\alpha)^{-1}x^{1-\alpha} \tag{6}$$

In this paper with focus on linear costs (a linearisation of (5)) and max-min as fairness criterion (6) ($w_d = 1$, $\alpha \rightarrow +\infty$). The choice of one particular fairness criterion is not of great importance in this context. We adopt max-min and linearised costs for being able to solve very large problems in realistic access networks. Here, we neglect all issues related to a real routing protocol implementing multi-path routing and flow control whilst focusing on the gain given by its deployment.

The optimal solution to (1) can be found by adopting an iterative LP formulation of the problem. At each iteration a linear sub-problem (whose formulation is omitted for the lack of space) is solved whom, at optimum, gives the same network *flow share* to every constrained demand: these former are then removed from the problem that iterates so on. Further details on this procedure can be found in [21] which is also similar to [22]. If there is enough resource for all demands flow control (however implemented) has no impact on the throughput and route selection is given by the optimum of the multi-commodity flow problem. In presence of congestion, flow control and routing jointly influence optimal route and bandwidth allocation.

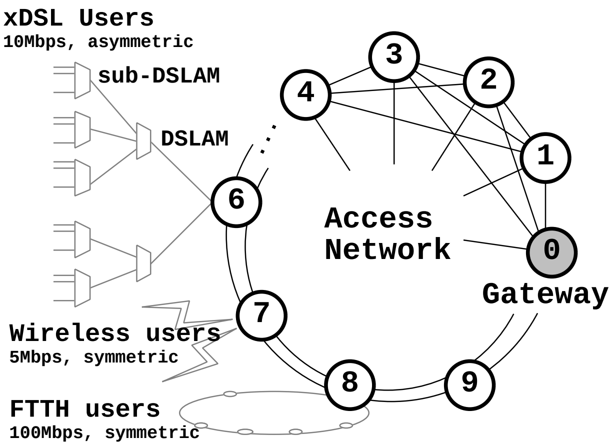

Fig. 1. Reference Network Scenario.

TABLE I
CURRENT AND FUTURE REFERENCE SCENARIOS

		Current	Future
Access Type	Internet gateway	0	0
	ADSL	1-6	1-3
	WiFi	7,8	4-6
	FTTH	9,10	7-10
Traffic Type	local (1-10)	25%	25%
	hot-spot (0)	75%	75%

ACCESS TECHNOLOGIES AND CAPACITIES

Network Segment	Link Technology	Capacity	
		Uplink	Downlink
Customer's Access	ADSL	1 Mbps	10 Mbps
	WiFi	5 Mbps	
	FTTH	100 Mbps	
Access	Ethernet, Optical	1–100 Gbps	

IV. REFERENCE NETWORK AND TRAFFIC SCENARIOS

We illustrate the reference network and traffic scenarios with the help of Fig. 1.

As we previously mentioned, we model an heterogeneous population of users that access to the network via different technologies. Specifically, we consider xDSL, Wireless and FTTH access to which correspond different *uplink* rates of 1 Mbps, 5 Mbps and 100 Mbps respectively. Moreover, FTTH and Wireless users have fully symmetric access, whereas xDSL ones have asymmetric downlink rates of 10 Mbps.

As represented in Fig. 1, behind each access node we consider an homogeneous population of 1000 users, so that heterogeneity of access technology will emerge when considering the whole network.

We consider an access network consisting of $N=$11 nodes; we vary the node degree L, exploring all possible topologies from the ring ($L=2$) to the full mesh ($L=10$). Unless otherwise stated, however, we consider the access network to be a partial regular mesh, where nodes have degree $L=6$ and links are bidirectional. We suppose all links have the same capacity C which varies in the 1–100 Gbps to take into account current metropolitan link capacities and their possible future upgrades.
Nodes are labeled from 0 to 10, with node 0 (light-gray shaded) acting as a gateway towards the big Internet. Access nodes are randomly assigned to access technologies according to scenarios described in Tab. I.

Nowadays, most users have xDSL access technologies while FTTH and Wireless are exploited by a smaller number of people: however, FTTH subscribers are growing extremely fast. At the same time, in the future we will likely assist to an *increased heterogeneity* of access types, e.g., due to the deployment of new technologies, such as WiMAX/LTE. However, it is unlike that a single technology will entirely take over the others; instead, it is more reasonable to envision that users will likely use different technologies depending on their location, and even use several technologies at the same time.

This is reflected in Tab. I, which reports several details concerning the current and future scenarios used as reference for performance evaluation in Sec. V-A and Sec. V-B respectively. At the same time, since we have no means to be sure that the future will look alike to the one we depicted so far, we prefer to investigate a wider range of access technology breakdown in Sec. V-C, so to ensure that our conclusions will be valid to a more general extent.

As concerns the traffic matrix, we suppose a given percentage of traffic demands is directed to gateway, whereas other demands remain local to the access network (nodes 1-10). In the latter, destinations are chosen uniformly at random, and in both cases users generate traffic proportionally to their available bandwidth.

Nowadays the most significant portion of traffic is directed towards the big Internet, whereas, as early anticipated, *two contrasting forces* may shape the future scenario. We may observe either an increase of local traffic (e.g., as a result of the adoption of latency-aware [6], [7] peer selection mechanisms) or an increase of remote traffic (e.g., due to the Internet video traffic popularity phenomenon [5]).

As such, we will explore a wide range of Internet Hot-spot ratios H, namely spanning from 5% to 100%. Notice that we explicitly take into account an exhaustive interval so to have a full-blown view of the multipath performance bounds.

We aggregate altogether all demands coming from users behind a given access node and targeting users behind the same access node: in this way, we focus on the transport of aggregated demands on the access network.

Traffic aggregation is also the only possibility at this stage as any realistic implementation would require to scale up to a very large number of flows in progress.

V. PERFORMANCE EVALUATION

This section compares the performance of Multi-Path (MP) versus Min-Cost (MC) routing strategies. We report the numerical results of the iterative LP problem formulated in

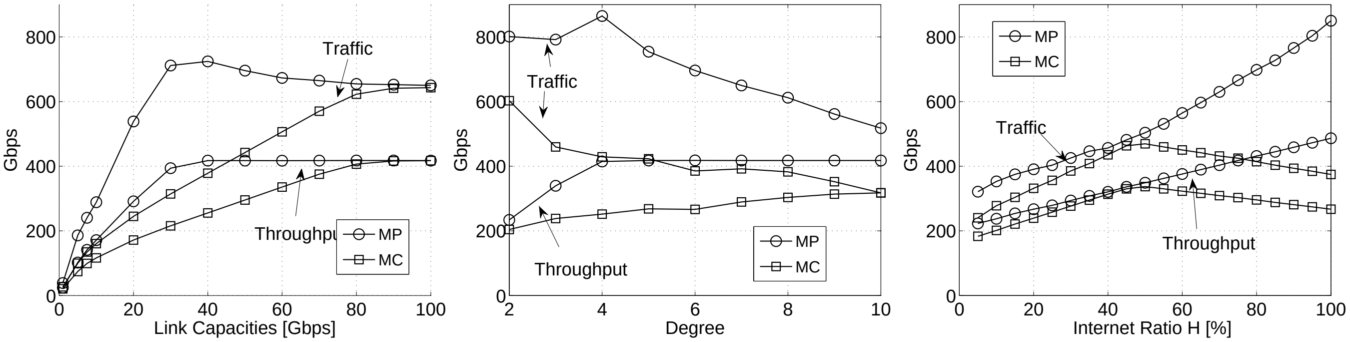

(a) (b) (c)

Fig. 2. Current scenario: network traffic for varying (a) link capacity (b) node degree and (c) traffic locality ratio.

Sec. III applied to the current and future reference scenarios in Sec. V-A and Sec. V-B respectively, widening furthermore the boundary of the possible scenarios in Sec. V-C.

For each considered scenario, results are averaged over 10 different instances of the problem, where we fix a traffic matrix and a topology and we randomize the type of access associated to each node.

As performance measures, we consider both network-centric and user-centric indexes. For instance, network performances will be expressed in terms of the *net* amount of traffic demands that the access network is able to serve, as well as in terms of the *total* amount of traffic flowing into the network. We instead express user-centric performance in terms of user *satisfaction*, defined as the bandwidth allocated to a demand normalized over its requested rate: normalization hides how this satisfaction is distributed among users with different access technologies and has also no dependency on the volume of the transported demand. We avoid to consider other QoS metrics, e.g. delays, because our focus is on the Internet best effort network. In the context of diffserv networks with specific QoS constraints the objective may be modified and adapted ad hoc but this is not the goal of this paper.

A. Current scenario

We start by comparing the performance of MC and MP in the current network scenario. Fig. 2 reports, in three different plots, several performance measures that help understanding the network state and user satisfaction from different perspectives. Let us start by adopting a network-centric perspective and evaluate the network efficiency in servicing the traffic demands. The net amount of demands injected into the network as well as the total amount of data needed to dispatch them are depicted in Fig. 2-a as a function of the individual link capacities, varying from 1 to 100 Gbps. First of all, notice that for any given link capacity, MP is able to serve a significantly higher amount of net demands with respect to MC.

Then, it is important to notice that the total amount of network traffic used by multi-path routing is *not monotone* with the link capacities, but rather increases steadily if needed to satisfy any demand. To see why this happens, consider first the case in which link capacities are very large, i.e., right portion of Fig. 2-a: multi-path routes serves every demand as min-cost. Thus, for 100 Gbps links, there is no difference between the amount of traffic MP and MC are able to serve. However, if we reduce link capacities, multi-path starts to route demands in alternative paths, in order to keep users satisfaction at the same level. Thus, MP uses in this case alternative paths, longer with respect to the shortest path used by MC, and consequently the total network traffic increases. If we further decrease link capacities, we eventually reach a point where the available network capacity becomes scarce and users satisfaction is forced to diminish: at this stage, it is unavoidable that network traffic decreases as well, which explain the behaviour observed in of Fig. 2-a.

Tab. II, reports the average satisfaction of demands with respect to their rates. Demands are classified with respect to their requested demand whether smaller than 1Gbps (Class 1), between 1Gbps and 10Gbps (Class 2) and larger that 10Gbps. It can be observed that, at optimum, only high rate demands get the benefits brought by path diversity. This means that customers with high rate lines would find a bottleneck at the access network unless path diversity is exploited.

Then, let us compare MC and MP performance for varying access network topologies when individual link capacities is set to 50 Gbps. More precisely, Fig. 2-b depicts the overall traffic and throughput as a function of the node degree L, where $L = 2$ implies that topology is a ring and $L = 10$ a full mesh. It can be seen that, as long as few “short-cuts” are added to the ring, MP is able to take profit of them so to serve an increased amount of traffic. When the degree grows beyond $L > 4$, there is no further advantage in the amount of traffic MP is able to serve – i.e., MP throughput saturates. Yet, adding more links obviously shorten the length of MP paths, which explains why the total traffic offered to the network continues to decrease as L increases.

Finally, considering a partial mesh $L = 6$ with individual links of $C = 50\,Gbps$, we vary the amount of traffic directed toward the Internet hot-spot H, and report results in Fig. 2-

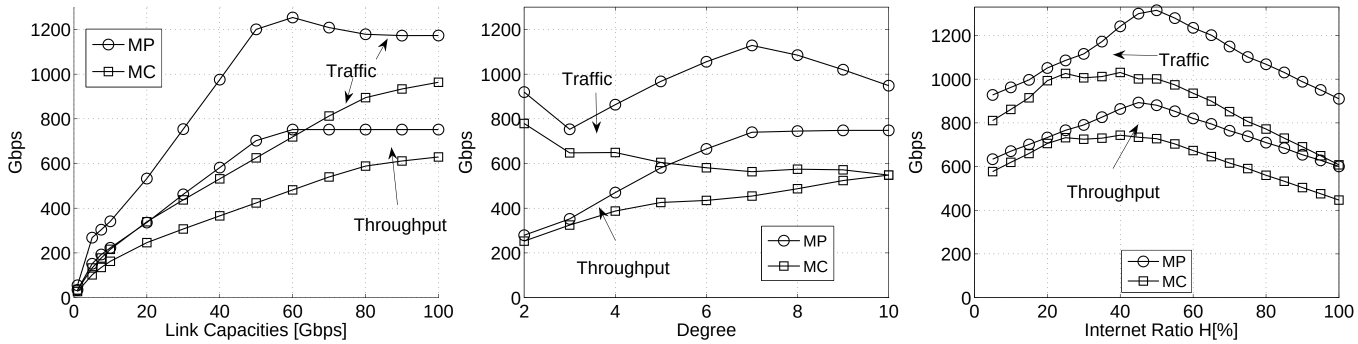

(a) (b) (c)

Fig. 3. Future scenario: network traffic for varying (a) link capacity (b) node degree and (c) traffic locality ratio.

AVERAGE SATISFACTION PER CLASS

		Class 1	Class 2	Class 3
C=1	MC	93%	24%	2%
	MP	95%	27%	4%
C=10	MC	100%	91%	25%
	MP	100%	96%	46%
C=50	MC	100%	100%	73%
	MP	100%	100%	100%

TABLE II
CURRENT SCENARIO

		Class 1	Class 2	Class 3
C=1	MC	98%	16%	2%
	MP	100%	19%	2%
C=10	MC	100%	90%	23%
	MP	100%	100%	40%
C=50	MC	100%	100%	86%
	MP	100%	100%	94%

TABLE III
FUTURE SCENARIO

c. MC curves reach a maximum around $H = 50\%$; at this point the average satisfaction is at maximum (about 100%): all demands are satisfied as there are no bottlenecks. Exceeding $H = 50\%$, as there are two FTTH access nodes, their total demand exceeds 100 Gbps which is their maximum bandwidth towards the gateway. Beyond this operational point, the aggregation network becomes a bottleneck, thus the traffic demand cannot be fully routed toward the gateway and the overall throughput decreases. This phenomenon is instead absent in the case of MP routing, where the use of multiple paths is able to sustain any admissible traffic by routing demands on less congested paths. MP avoids network bottlenecks, as the maximum requested demand remains below 600 Gbps ($L \times C$), being the maximum capacity MP can route from/towards the gateway.

B. Future scenario

Let us now consider the evolutionary scenario, whose main peculiarity is an increased access heterogeneity. For the moment, however, we consider that the locality ratio is *unchanged* with respect to the previous scenario, so that we can evaluate the impact of the access technology in isolation.

Fig. 3-a again depicts the net amount of demands injected into the network as well as the total amount of data needed to dispatch them as function of link capacities. Results allow us to conclude that MP needs 60 Gbps links satisfy every demand while single-path routing is not able to satisfy all requests, and the network would require far more than 100 Gbps links. The picture also shows that MC would need about 50% additional capacity in order to transfer the same amount of data of MP.

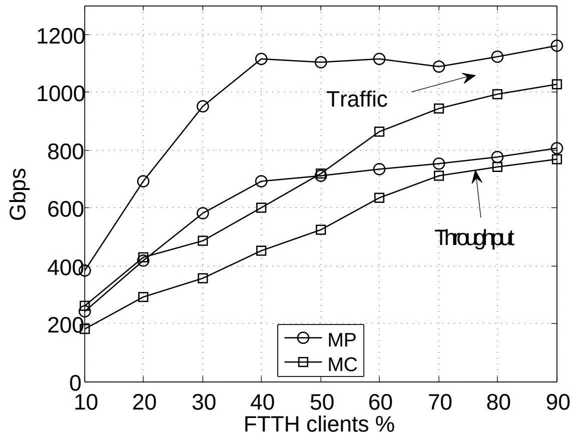

Fig. 4. Performance as a function of "access diversity".

As for the current scenario, Tab. II confirms only high rate demands can increase their throughput thanks to path-diversity.

Fig. 3-b shows MP requires a degree $L > 6$ in order to satisfy all demands while MC is not able to satisfy them all even in a full-mesh topology.

Finally, Fig. 3-c considers the performance for an access network with degree $L = 6$ and $C = 50 Gbps$ as a function of the traffic directed towards the Internet gateway H.

With respect to the "current scenario" the demand towards the gateway, exceeds 600 Gbps ($L \times C$) so that even MP does not attain 100% satisfaction when $H > 50\%$. However satisfaction is close to 100%. Beyond this operational point the throughput of both routing schemes, MP and MC, decreases as they all attain the network bottleneck which becomes the gateway when $H = 100\%$.

C. Access heterogeneity

In the presented scenarios we have seen that routing performance are heavily related to the distribution of the access rates. In this section we go through this observation in order to deeply understand the phenomenon.

We fix link capacities to C=50 Gbps, node degrees to L=6, traffic locality ratio to H=75% and variate the distribution of the access technology. We divide users into two categories: high (FTTH) and low rates (WiFi and ADSL). In more detail, we vary the percentage of FTTH users from 10% to 90%, and assume that the remaining portion of access can be either WiFi or ADSL. For instance, when a single node (or, X nodes) is FTTH, we partition at random the remaining 9 (or, $10 - X$) nodes between WiFi and ADSL access, and repeat each experiment 10 times varying the traffic matrix and WiFi/ADSL breakdown.

Fig. 4 shows that the more the access diversity the more MP gains with respect to MC. When only one access node connects FTTH clients (left hand side of Fig. 4) the gateway is the unique hot-spot, and the gain is limited. As more access nodes connect FTTH clients there are more local exchanges (e.g. local high rate seeders in P2P swarms) and the maximum gain is reached when the FTTH ratio attains 40%. Indeed at this point the traffic matrix is very unbalanced because multiple hot-spots show up. However, as the main part of the clients use FTTH the traffic matrix is far more uniform being the gateway the unique hot-spot (right hand side of Fig. 4).

VI. Discussion and Conclusion

This work focused on the design and analysis of architectures and routing mechanisms in next generation metro access networks. We considered a framework of optimal routing and bandwidth allocation, which is able to take into account peculiar user characteristics such as access capacities.

By means of numerical evaluation, we assess the gain that can be obtained exploiting path diversity for large networks with a large number of demands. More specifically, we have investigated the impact that the traffic locality and the aggregation topology have in determining the performance of the network metro segment. Multi-path is able to more efficiently exploit additional available capacity to fulfil additional demands, and better use existing resources, potentially avoiding the need for costly links capacity upgrades.

Multi-path routing reduces the extent of the link capacity upgrade by almost a factor of two with respect to min-cost routing – which confirms it to be an appealing strategy for both current and future network architectures.

Our study suggests that next generation networks, in the short/medium term, should evolve to more meshed topologies in order to exploit path diversity and implement multi-path routing strategies, e.g. better exploiting MPLS and flow control mechanisms. In the long term, in a clean-slate perspective, the access network should definitely implement native multi-path routing and flow control primitives. The control plane should implement agile route ranking to monitor network state and to properly balance network demands in presence of overload.

Acknowledgement

The industrial partner has been partially funded by EU FP7 Integrated Project 4WARD The academic partner has been funded by Celtic TIGER, a project of the Eureka cluster.

References

[1] eMule, http://www.emule-project.net/

[2] Bit Torrent, http://www.bittorrent.com

[3] PPlive, http://www.pplive.com

[4] SOPCast, http://www.sopcast.com

[5] "Cisco Visual Networking Index – Forecast and Methodology, 2007–2012", available at http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360.html

[6] S. Ratnasamy, M. Handley, R. Karp, S. Shenker "Topologically-Aware Overlay Construction and Server Selection" *IEEE Infocom'02*, 2002

[7] F. Dabek, J. Li, E. Sit, J. Robertson, M. Kaashoek, R. Morris", "Designing a DHT for low latency and high throughput", *USENIX Symposium on Networked Systems Design and Implementation (NSDI'04)*, March 2004.

[8] D.G. Cantor and M. Gerla M., "Optimal routing in a packet switched computer network," *IEEE Transaction on Communications*, Vol. C, No. 23, pp. 1062-1069, Oct. 1974.

[9] R. Gallager, "A Minimum Delay Routing Algorithm Using Distributed Computation," *IEEE Transaction on Communications*, Vol. 25, No. 1, Jan. 1977.

[10] S. J. Golestaani. "A Unified Theory of Flow Control and Routing in Data Communication Networks," *Ph.D. Thesis, MIT Cambridge Lab for Information and decision systems*, , December 1979.

[11] L. E. Hahne "Round Robin Scheduling for Fair Flow Control in Data Communication Networks," *Ph.D. Thesis, MIT Cambridge Lab for Information and decision systems*, , December 1979.

[12] H. Han, S. Shakkottai, C.V. Hollot, R. Srikant, D. Towsley, "Multipath TCP: a joint congestion control and routing scheme to exploit path diversity in the internet" *IEEE/ACM Transaction on Networking*, Vol. 14 , No. 6, Dec. 2006

[13] J. He, M. Bresler, M. Chiang, J. Rexford, "Towards Robust Multi-Layer Traffic Engineering: Optimization of Congestion Control and Routing," *IEEE Journal on Selected Areas in Communications*, Vol. 25, No. 5, June 2007

[14] S. Kandula, D. Katabi, B. Davie, A. Charnie "Walking the tightrope: responsive yet stable traffic engineering," *ACM SIGCOMM 2005*

[15] F. Kelly, T. Voice, "Stability of end-to-end algorithms for joint routing and rate control," *ACM SIGCOMM Computer Communication Review*, 35:2, pp. 5-12, 2005

[16] T. Voice, "Stability of Multi-Path Dual Congestion Control Algorithms," *to appear in IEEE/ACM Transactions on Networking*

[17] X. Lin, N. B. Shroff, "Utility maximization for communication networks with multipath routing," *IEEE Transactions on Automatic Control*, May 2006. Vol. 51 No. 5

[18] P. Key, L. Massoulié, D. Towsley, "Path selection and multipath congestion control," *IEEE Infocom'07*

[19] S. Oueslati, J. Roberts "Comparing Flow-Aware and Flow-Oblivious Adaptive Routing," *Proceedings of CISS'06*

[20] F. Kelly, "Charging and rate control for elastic traffic," *European Transactions on Telecommunications*, Vol. 8, pp. 33–37, 1997

[21] L. Muscariello and D. Perino, "Early Experiences in Traffic Engineering Exploiting Path Diversity: A Practical Approach," *INRIA Technical Report No. 6474*, 2008

[22] Nace, D; Linh Doan, N; Gourdin, E; Liau, B; Computing optimal max-min fair resource allocation for elastic flows. IEEE/ACM Transactions on Networking 16(6): 1272-1281 (2006)